arXiv:1607.04212v1 [cond-mat.str-el] 14 Jul 2016

A comparison between methods of analytical continuation for bosonic functions

J. Schött,[1] E. G. C. P. van Loon,[2] I. L. M. Locht,[1,2] M. I. Katsnelson,[2] and I. Di Marco[1]

[1] *Dept. of Physics and Astronomy, Uppsala University, Box 516, SE-75120 Uppsala, Sweden*
[2] *Institute of Molecules and Materials, Radboud University of Nijmegen, Heyendaalseweg 135, 6525 AJ Nijmegen, The Netherlands*

In this article we perform a critical assessment of different known methods for the analytical continuation of bosonic functions, namely the maximum entropy method, the non-negative least-square method, the non-negative Tikhonov method, the Padé approximant method, and a stochastic sampling method. Three functions of different shape are investigated, corresponding to three physically relevant scenarios. They include a simple two-pole model function and two flavours of the non-interacting Hubbard model on a square lattice, i.e. a single-orbital metallic system and a two-orbitals insulating system. The effect of numerical noise in the input data on the analytical continuation is discussed in detail. Overall, the stochastic method by Mishchenko *et al.* [Phys. Rev. B **62**, 6317 (2000)] is shown to be the most reliable tool for input data whose numerical precision is not known. For high precision input data, this approach is slightly outperformed by the Padé approximant method, which combines a good resolution power with a good numerical stability. Although none of the methods retrieves all features in the spectra in the presence of noise, our analysis provides a useful guideline for obtaining reliable information of the spectral function in cases of practical interest.

PACS numbers: 71.10.Fd, 71.15.Dx, 02.70.Hm

I. INTRODUCTION

Strongly correlated materials exhibit a wide range of exotic physical phenomena, ranging from magnetism to superconductivity [1]. This rich physics is interesting for the scientific community since it has a high potential for future technological applications. In the last decades several theories have been developed to describe the electronic structure of strongly correlated materials with a good accuracy, e.g. the combination of density-functional theory (DFT) and dynamical mean-field theory (DMFT) [2–5]. Practical calculations for systems at finite temperature are usually performed using the Green's function formalism for complex energies [6, 7]. In this approach determining physical observables requires an analytical continuation from complex energies (Matsubara frequencies) to real energies, as depicted in Fig. 1.

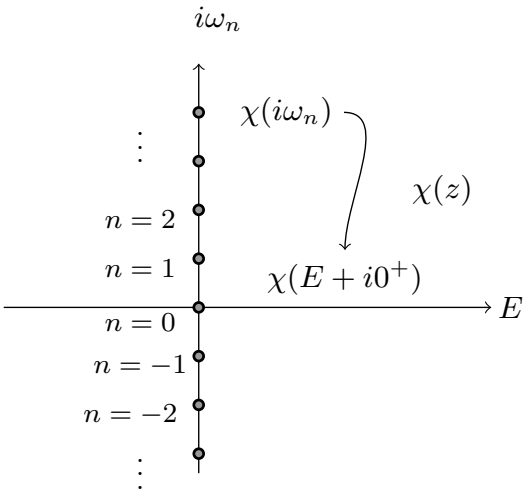

Figure 1. Schematic illustration of the analytic continuation of a bosonic response function $\chi(z)$. The values $\chi(i\omega_n)$ at the Matsubara frequencies $i\omega_n$ are used to reconstruct the function $\chi(E+i\delta^+)$ just above the axis of real energies.

Traditionally, the analytical continuation of one-particle Green's functions has received the most attention, since it is related to the one-particle spectral function. The latter is not only of fundamental importance, but can also be probed (more or less) directly via various types of photoemission experiments. Over the last few years however, dynamical *two-particle* quantities became more and more important [8–16]. Significant examples include the electron energy loss spectrum (EELS) [9, 13] and dynamical susceptibilities [15, 17, 18]. Physically, the two-particle spectrum contains several interesting features, such as low-energy dispersive Goldstone modes, Landau damping of collective modes and sharp transitions between isolated energy levels. Even in a weakly interacting Fermi liquid, the two-particle spectrum presents some of these interesting features [7]. From a more technical point of view, two-particle quantities also play an important role in the treatment of nonlocal interaction in DMFT-based approaches such as extended DMFT (EDMFT) [19–23] and its extension dual boson [24], by describing the feedback of collective excitations on the one-particle spectrum.

Similarly to the one-particle Green's function, two-particle quantities are usually obtained on the Matsubara axis. Some physical information, such as the occurrence of charge-order transitions [11, 25, 26], can be obtained directly from computational data on the Matsubara axis at zero-frequency. However, the full dynamical susceptibilities, as well as related observables such as the EELS and the plasmon spectrum, require an analytical continuation to real energies. It is easy to understand why

finding a reliable method for the analytical continuation of two-particle functions has evolved from a niche problem to a necessity in computational many-body physics.

The usual methods of analytical continuation of fermionic functions can, after some manipulation, also be used for bosonic functions [27]. It was previously shown that the maximum entropy method (MEM) is able to perform the analytical continuation of the optical conductivity with a good accuracy [28]. More recently, in the context of EDMFT, Huang *et al.* [12] focused on the retarted interaction, emphasizing that a modified version of MEM [29] leads to a good analytical continuation. While these studies provide a very interesting and useful analysis, they do not cover the applicability of various methods to functions of different shapes and characters. So far, a more complete study of the analytical continuation problem for bosonic functions is lacking. In this work we intend to provide such an analysis, by considering five different methods of analytical continuation in three physically relevant scenarios. In addition to the MEM [27, 28, 30–35], we will also test the non-negative least-square (NNLS) method [36], the non-negative Tikhonov (NNT) method [37], the Padé approximant (Padé) method [38, 39], and Mishchenko's stochastic sampling method [40, 41]. Furthermore, we will investigate the effect of numerical noise on the quality of the continuation, in connection to the usage of quantum Monte Carlo techniques (see e.g. Ref. 42).

This article is structured as follows; in Sec. II we review bosonic Green's functions and their analytical and symmetry properties. In Sec. III a brief introduction to the various continuation schemes is given. Then, in Sec. IV, we describe the different testing cases and highlight the reasons why they are interesting. In Sec. V, we show and analyze the results of the tests. Finally in Sec. VI we present the conclusions of our study, including a comparative summary of the performance of the continuation methods investigated.

II. BOSONIC GREEN'S FUNCTIONS

The analytical continuation consists in obtaining a function $\chi(z)$ in the whole complex plane from a finite set of values. The focus in this work is on functions with bosonic symmetry that are known on a finite number of bosonic Matsubara frequencies $i\omega_n = 2n\pi Ti$, where T is the temperature and n is an integer number. From this partial knowledge, we intend to extract the values of the function in the entire complex plane, and especially just above the real axis, for $z = E + i\delta$.

The analytical continuation problem can be formulated in terms of the spectral function $\rho(E)$ by means of the Hilbert transform

$$\chi(i\omega_n) = \int_{-\infty}^{\infty} dE \frac{1}{i\omega_n - E}\ \rho(E)\ . \tag{1}$$

For bosonic functions, the spectral function is an odd function, i.e. $\rho(E) = -\rho(-E)$, which allows us to simplify Eq. (1) to

$$\chi(i\omega_n) = \int_{0}^{\infty} dE \frac{-2E}{\omega_n^2 + E^2}\ \rho(E)\ . \tag{2}$$

From this equation we find that for large Matsubara frequencies $\chi(i\omega_n)$ is asymptotically proportional to ω_n^{-2}. Without the odd symmetry of ρ, one would obtain an asymptotic behavior as ω_n^{-1}, like for the one-particle Green's function. Furthermore, $\chi(i\omega_n)$ is purely real and $\chi(i\omega_n) = \chi(-i\omega_n)$. It follows from the analytical properties of $\chi(z)$ that $\chi(i\omega_n)$ decreases monotonously as a function of $i\omega_n$ [43]. The correct symmetry can be enforced in the analytical continuation procedure either explicitly by using Eq. (2) or implicitly by including both positive and negative Matsubara frequencies as input points. This issue is studied in more detail in App. A.

III. CONTINUATION ALGORITHMS

The analytical continuations are performed by means of five different methods. We use our own, in-house implementations for all algorithms except for Mishchenko's method. For Mishchenko's method we use an MPI-parallelized version of the original code presented in Ref. 40. Here, a brief introduction to these algorithms is given.

1. *NNLS method*

Discretizing the integral in Eq. 2 makes it possible to reformulate the problem as a system of linear equations

$$\chi(i\omega_n) = \sum_j \underbrace{f_j \frac{-2E_j}{\omega_n^2 + E_j^2}}_{K_{n,j}}\ \rho_j\ , \tag{3}$$

where f_j is a quadrature weight and $K_{n,j}$ is the matrix to invert. However,this approach would not work in practice, since obtaining $\rho(E)$, once $\chi(i\omega_n)$ is given, is an ill-posed problem. The NNLS method solves Eq. (3) in a least-square sense and stabilises the solution by enforcing the known symmetry property $\rho(E)E \geq 0$. This means finding a non-negative solution to a least-square problem:

$$\min_{\boldsymbol{\rho}\geq \mathbf{0}} \|\boldsymbol{\chi} - K\,\boldsymbol{\rho}\|^2\ . \tag{4}$$

Here, $\boldsymbol{\chi}$ and $\boldsymbol{\rho}$ are vectors containing the values $\chi(i\omega_n)$ and ρ_j on the discrete points of Equation (3). The NNLS problem can be solved iteratively as described extensively in Ref. 36.

2. *NNT method*

For ill-posed problems, regularisations are commonly used. One of the most famous regularisation methods

is the Tikhonov method [37]. Applying the Tikhonov method to the problem of analytical continuation defines the NNT method as

$$\min_{\boldsymbol{\rho}\geq\mathbf{0}} \|\boldsymbol{\chi} - K\,\boldsymbol{\rho}\|^2 + \alpha\,\|\boldsymbol{\rho}\|^2\,, \tag{5}$$

where α is a weight parameter that can be determined by the L-curve method [37] which selects α corresponding to the smallest value of $\ln(\|\boldsymbol{\chi} - K\,\boldsymbol{\rho}_\alpha\|^2) + \ln(\|\boldsymbol{\rho}_\alpha\|^2)$. Here $\boldsymbol{\rho}_\alpha$ denotes the solution of Eq. 5 for a fixed α.

3. MEM

Another famous regularisation method is the Maximum Entropy Method (MEM) [27, 28, 30–35]. This method maximises the Neumann entropy for the spectral function. Formulated as a minimisation problem, the equations become:

$$\min_{\boldsymbol{\rho}\geq\mathbf{0}} \|\boldsymbol{\chi} - K\,\boldsymbol{\rho}\|^2 + \alpha\,\|S[\boldsymbol{\rho}]\|\,, \tag{6}$$

where the entropy

$$\begin{aligned} S[\rho] &= \int dE\rho(E)\ln\left(\frac{\rho(E)}{m(E)}\right) \\ &\approx \sum_j f_j\rho_j \ln\left(\frac{\rho_j}{m_j}\right) \end{aligned} \tag{7}$$

contains a default model $m(E)$, which can incorporate a priori knowledge about the spectral function.

4. Padé approximant method

The Padé approximant method [38, 39] is based on a fitting procedure. This method starts with a rational polynomial ansatz for $\chi(z)$ with unknown coefficients, and fits this ansatz to the known values at the Matsubara axis $\chi(i\omega_n)$. In this way one can reconstruct the function in the entire complex plane. The Padé method can be formulated as a matrix problem [38], which improves its numerical accuracy. As recently proposed [44], one can enhance the stability of the continuations (especially in presence of Matsubara noise) by taking an average of several Padé approximants, obtained by varying the number of Padé coefficients and Matsubara points in the fitting procedure.

5. Mishchenko's stochastic sampling method

Finally, we consider the stochastic sampling method proposed by Mishchenko *et al.* [40, 41]. In this method the spectral function is approximated by a set c of rectangles R, i.e.:

$$\rho_c(\omega) = \sum_{R\in c} R_{\{h,b,m\}}(\omega)\,.$$

The rectangles R are defined by their height h, width b and center position m. A stochastic algorithm updates the rectangles randomly and accepts the changes according to a Metropolis algorithm based on the difference between the known function $\chi(i\omega_n)$ and the corresponding function obtained from $\rho_c(\omega)$. Combining this update scheme with a deterministic minimization of the Matsubara difference further improves the performance of the method. An average over different independent Monte Carlo chains is taken to minimize the influence of noise. Since the algorithm uses a condition for the sum of the rectangle weights, we scale ρ by introducing $\tilde{\rho}(E) = \frac{-2}{E\chi_0}\rho(E)$. Instead of finding ρ by solving Eq. 2, we calculate $\tilde{\rho}$, which has to obey

$$\tilde{\chi}(i\omega_n) = \frac{\chi(i\omega_n)}{\chi_0} = \int_0^\infty dE \frac{E^2}{\omega_n^2 + E^2}\;\tilde{\rho}(E)\,. \tag{8}$$

For $n = 0$ we have the desired normalization condition $1 = \int_0^\infty dE\;\tilde{\rho}(E)$. Once $\tilde{\rho}$ is found we scale back to obtain ρ. This rescaling can in principle be used for the other methods as well, with the exception of Padé where the continuation is not done by solving Eq. 2.

IV. MODEL DESCRIPTIONS

To investigate the quality of the analytical continuation provided by the different methods we focus on test functions that are known in the entire complex plane, including the real energy axis. The functions are first evaluated at the Matsubara energies and then analytically continued to the real axis. In an ideal situation, the solution obtained by analytical continuation coincide with the original exact values. However, in practice, some differences arise due to the limited precision and amount of input data.

The continuation becomes more difficult when we add numerical noise to the Matsubara data, which models what happens in real calculations as by, e.g., Monte Carlo methods. There, the noise typically scales inversely with the square root of the computational time. As a practical example, the numerical noise in dual boson calculations [25] is often in the 0.1% range. The numerical noise is modelled by adding relative Gaussian noise to each Matsubara point, $\chi(i\omega_n)(1+\epsilon)$, where ϵ is sampled from a Gaussian distribution with zero mean and a standard deviation σ. We use the same σ for all frequencies. In realistic Monte Carlo simulations, the noise level will usually depend on the Matsubara frequency $i\omega_n$ in a complicated way.

It is important to individuate test functions that are realistic and cover bosonic functions of different character as one would encounter in practice. Naturally, the exact real energy solutions of many of the most interesting systems are not known, so we need to find realistic but tractable cases. Here, we consider three scenarios: (i) a two-pole model, (ii) the momentum resolved random phase approximation (RPA) solution of the 2d doped

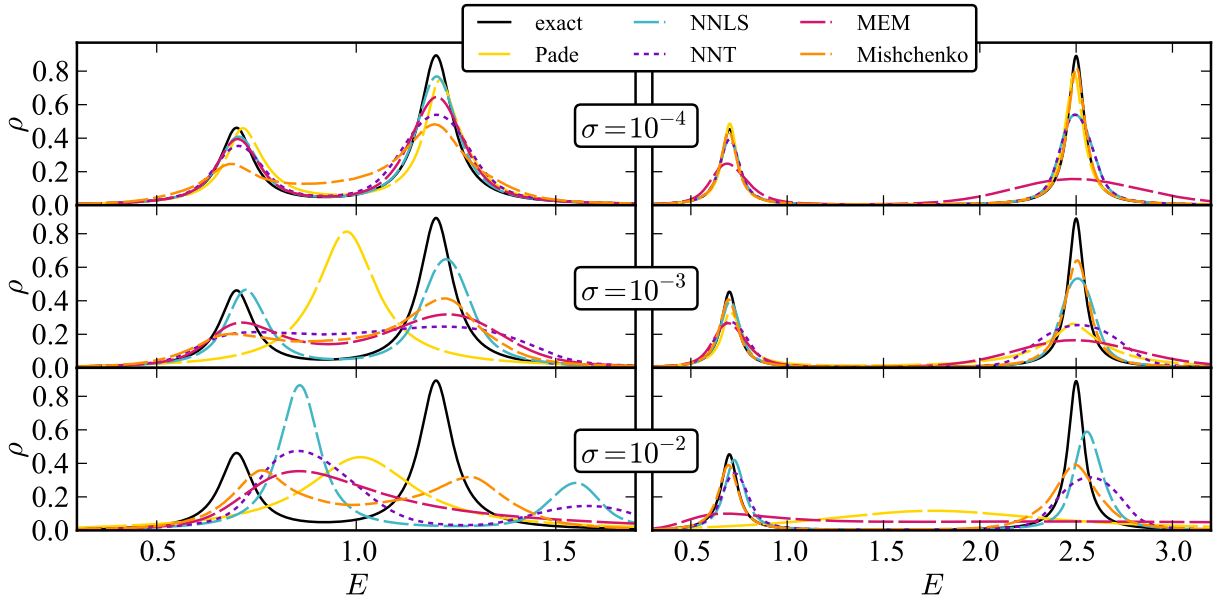

Figure 2. Spectra for the two-pole test function. Left panel $a_1 = 0.1$, $a_2 = 0.335663$, $E_1 = 0.7$ and $E_2 = 1.2$, right panel $a_1 = 0.1$, $a_2 = 0.7$, $E_1 = 0.7$ and $E_2 = 2.5$, as defined in Eq. 9.

non-interacting Hubbard model and (iii) that of a band insulator (with two identical bands separated by a gap). These three functions, which are explicitly described in Sec. V, cover sharp features, a sharp low-energy zero sound mode (that gets broadened by Landau damping at larger momenta), high frequency modes, features far away from zero energy and a combination of sharp features with a broad base.

All these functions are evaluated for Matsubara frequencies corresponding to an inverse temperature $\beta = 50$ and at a distance above the real axis $\delta = 0.05$. For all methods we used 100 Matsubara frequencies, except for Padé where a spectral average is done by varying the number of Matsubara points between 50 to 100 in steps of 4. It is important to use a sufficiently large number of Matsubara frequencies. We study the dependence of the quality of the continuation on the number of Matsubara frequencies in Sec. V A. We also investigate the effect of including a few negative frequencies in Appendix A.

V. RESULTS

In this section we present a comparison between the different methods of analytical continuation when applied to three selected test functions.

A. Two-pole model

The first function we address is a simple model function whose analytical structure is

$$\chi(z) = \frac{a_1}{z^2 - E_1^2} + \frac{a_2}{z^2 - E_2^2}\,, \tag{9}$$

where the parameters have the following values $a_1 = 0.1$, $a_2 = 0.7$, $E_1 = 0.7$ and $E_2 = 2.5$. For $\delta \to 0^+$, the exact spectrum has two distinct peaks at energies E_1 and E_2. Therefore this function has similarities to the Gaussian two-peak structure used in Ref. 12.

To understand how the position of the high energy peak influences the quality of the attained spectra, we also perform tests with parameters $E_2 = 1.2$ and $a_2 = 0.335663$. This choice preserves the height of both peaks in the two chosen setups. The two spectra, corresponding to $E_2 = 1.2$ and $E_2 = 2.5$, present different degrees of difficulty. For $E_2 = 1.2$, the two peaks in the spectrum are close to each other and therefore it is difficult to resolve them as two separate peaks. For $E_2 = 2.5$, on the other hand, the second peak is far away from the first peak, but also far away from zero energy (the complex $i\omega_n$-axis). This makes it difficult to resolve the distant peak.

In Fig. 2 the two exact spectra are reported, and compared to the analytical continuation obtained through the five methods for three different noise levels. It is known that the Padé method is much more sensitive to the precision of the input data than the other methods. The various noise levels are therefore chosen to illustrate the critical input precision where the Padé method starts to fail. These noise levels are realistic for Monte Carlo simulations [42].

As the top panels of Fig. 2 show, for relatively low noise ($\sigma = 10^{-4}$), all methods capture the two peaks and their positions correctly, for both sets of parameters. The Padé method performs the best, if we also consider the height of the peaks. Conversely, the MEM has the most difficulties to resolve the distant peak at $E_2 = 2.5$. This picture changes drastically for intermediate noise levels ($\sigma = 10^{-3}$), shown in the middle panels of Fig. 2. For $E_2 = 1.2$ (left panel), the Padé method is not able to reproduce the two peaks, and finds only the one at low energy. The other methods also gain spectral weight in between the peaks, but are still able to resolve both peaks. For $E_2 = 2.5$ (right panel), results are significantly better for all methods. Interestingly, the MEM performs worse than the Padé method. In situations where the input data is characterized by high noise ($\sigma = 10^{-2}$), depicted in the bottom panels of Fig. 2, both the MEM and the Padé method fail to find distinct peaks and merge them into one single broad peak. The other methods find distinct peaks, but locate them at the wrong energies, which is especially evident for $E_2 = 1.2$ (left panel). In any case, these methods are still able to describe the basic physics correctly, i.e. the presence of two well defined peaks. In general, Mishchenko's method seems to offer the best results across various levels of noise. Even for the highest noise ($\sigma = 10^{-2}$), Mishchenko's method gives an acceptable description of the exact spectra. Nevertheless one has to notice, if two peaks are close to each other, this approach tends to increase the spectral weight in between the peaks. As a matter of fact, this problem is not an intrinsic property of Mishchenko's algorithm, but is related to an insufficient number of local updates in the sampling. The dependence of Mishchenko's results with respect to the number of local updates is investigated in Appendix B. For the two-pole model, for small noise levels, Mishchenko's method can achieve results as good as Padé, although at the cost of a much bigger computational effort.

Finally, from the comparison of the left and right panels of Fig. 2, one can conclude that partially overlapping peaks seem to be more difficult to describe than having one peak close and one peak far from zero energy. This is not so trivial, since usually analytical continuation methods are supposed to perform well to describe (even complex) structures close to zero energy.

Number of Matsubara frequencies

It is also useful to illustrate how the quality of the analytical continuation depends on the number of input points, e.g., the number of Matsubara frequencies. The quality of the continuation can be measured through the difference between the obtained $\rho(E)$ and the exact function on the real axis $\rho_{\text{exact}}(E)$ in the following way

$$\text{Error} = \frac{\int_0^5 dE\, |\rho(E) - \rho_{\text{exact}}(E)|}{\int_0^5 dE\, |\rho_{\text{exact}}(E)|}\,, \qquad (10)$$

where the upper bound of the integrals is chosen to be larger than the actual extension of the function. The error measure given by Eq. (10) for the two-pole model with noise $\sigma = 10^{-4}$ is plotted in Fig. 3 with respect to different number of the maximum Matsubara index, $n_{\max}$. Note that for the Padé method the notion of $n_{\max}$ is a bit more complicated to define, since an average is done over several continuations with varying numbers of input points and Padé coefficients [44]. To make optimal usage of the averaging procedure, both the number of input points and Padé coefficients take values between $n_{\max} - 47$ and $n_{\max} + 1$ in steps of 4, under the constraint of not having more coefficients than Matsubara points. If $n_{\max} - 47 < 4$, then 4 is used as the lower boundary instead.

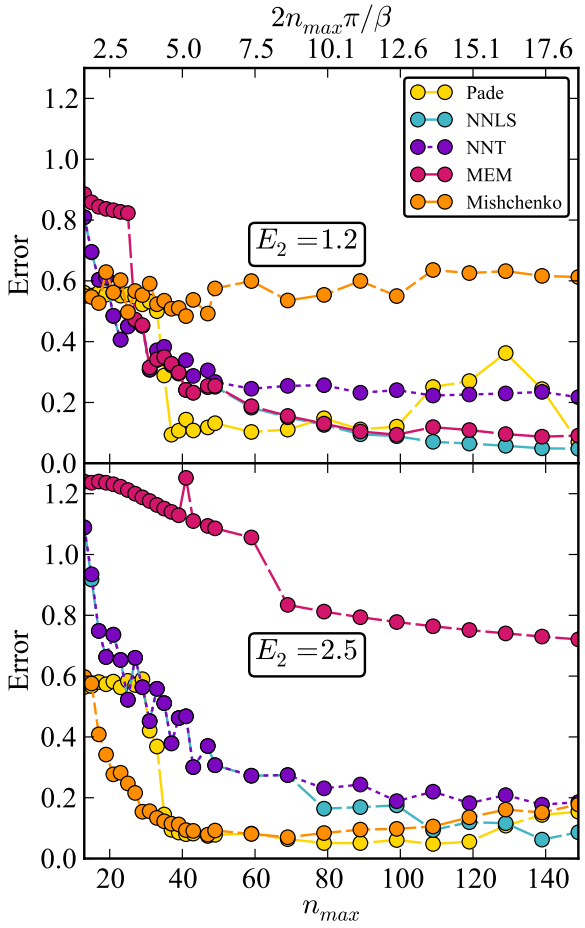

Figure 3. Integrated real axis error for the two-pole model as a function of the highest Matsubara point index $n_{\max}$ used. Upper panel $a_1 = 0.1$, $a_2 = 0.335663$, $E_1 = 0.7$ and $E_2 = 1.2$, lower panel $a_1 = 0.1$, $a_2 = 0.7$, $E_1 = 0.7$ and $E_2 = 2.5$, as defined in Eq. 9. Matsubara noise level $\sigma = 10^{-4}$ is used.

In Fig. 3, the general trend is that the quality of the continuation initially improves with an increasing number of Matsubara points. At some point, the error stabilises. When the typical energy scale on the real axis (E_2) increases, this stabilisation point moves to higher

Matsubara frequencies. However, there are small variations to this general picture for the different methods.

The error between the Padé continuation and the exact function quickly drops at $n_{\max} \approx 30$ and is very small until $n_{\max} \approx 100$ (or 120 for $E_2 = 2.5$). When taking into account more Matsubara points, the error increases. We attribute this behavior mainly to the number of physical continuations in the averaging procedure. For small $n_{\max}$, the total number of continuations is limited. For large $n_{\max}$, the number of physical continuations decreases significantly. For this particular test, at least 30 to 40 physical continuations where needed, to optimally use the power of the averaging procedure. The continuations obtained with NNT and NNLS, on the other hand, continue improving when additional Matsubara frequencies are taken into account. However the improvement becomes smaller and smaller. The MEM error quickly drops at $2n_{\max}\pi/\beta \approx 3E_2$ and decreases only slightly for higher $n_{\max}$. The error of the Mishchenko continuation stabilises at a relatively small number of Matsubara frequencies. For this method, we have used the same number of local updates (15000) for all continuations. The behaviour for more local updates is discussed in Appendix B. It is conceivable that adding additional frequencies does give extra information, but also requires a longer runtime of the stochastic sampling.

B. Doped Hubbard model

For the second test, we use a metallic system, namely the two-dimensional non-interacting Hubbard model on a square-lattice with nearest-neighbor interaction. The correlation function depends on the Matsubara frequency and on the momentum $\mathbf{q}$. According to RPA it is given by [7, 45]

$$\chi(i\omega_n, \mathbf{q}) = \sum_{\mathbf{p}} \frac{n_F(\epsilon_{\mathbf{p}}) - n_F(\epsilon_{\mathbf{p}+\mathbf{q}})}{\epsilon_{\mathbf{p}} - \epsilon_{\mathbf{p}+\mathbf{q}} + i\omega_n}, \tag{11}$$

where the sum denotes an average over the Brillouin zone, n_F is the Fermi distribution and the dispersion relation is

$$\epsilon_{\mathbf{p}} = -2t\left[\cos p_x + \cos p_y\right]. \tag{12}$$

The energy scale of this model is given by the half bandwidth $D = 4t$, where t is the hopping parameter. Therefore, we fix the energy unit to $D = 4t = 1$. The chemical potential is set to $\mu = -0.5$, which corresponds approximately to a filling of 0.185 per spin flavor.

In Fig. 4 the analytical continuation of the momentum resolved RPA susceptibility for all methods is reported. The panels in the first three columns correspond to data for various levels of noise σ. The panels in the last column illustrate the exact spectrum on the real axis, repeated in each row for an easier comparison.

At small wavevectors (close to Γ), the susceptibility exhibits a single low-energy mode, the zero sound mode. This mode consists of low energy excitations close to the Fermi surface. Further away from Γ, Landau damping broadens this zero sound mode. All methods correctly capture the zero sound mode. However, NNLS shows an unphysical splitting of this mode into two branches roughly halfway from $\Gamma \to X$. The tendency of NNLS to introduce spurious peaks is even more pronounced when going from $X \to M$, and in fact it is not even possible to illustrate all countour lines in Fig. 4 without compromising its readibility. The complex spectral evolution in the path $X \to M$ seems indeed rather difficult to describe for all methods, especially for high noise. At the M-point the exact spectrum acquires a seemingly simpler structure, consisting of a relatively flat bump between $E = 1$ and $E = 2$. This point is analysed more in detail in Fig. 5. It shows that it is not so easy to capture the broad flat feature instead of two separate peaks. Similarly to the previous test function, Padé correctly reproduces the spectrum at low noise, but gives a single sharp peak instead of a broad flat mode when σ is increased. NNLS performs very poorly here, with a spectrum of a few sharp peaks even at low noise $\sigma = 10^{-10}$. NNT does a lot better, compared to NNLS; the regularisation flattens the spectrum, although an unphysical local minimum remains at $\sigma = 10^{-4}$. Similarly, MEM also captures the general shape of the spectrum, but again an unphysical local minimum forms between the main peaks. Mishchenko's method is the only one that correctly reproduces the very flat spectrum between $E = 1$ and $E = 2$ even for high noise. These feature is consistent with the behavior observed in the top left panel of Fig. 2.

In general Mishchenko's method performs the best across various levels of noise, while NNLS leads to the poorest continuation for this function. As expected, the quality of the continuation obtained with Padé worsens quickly when σ is increased. Interestingly, NNT seems to perform slightly better for higher noise. For very low noise, the features are sharper and some spurious wiggles occur. These properties influence strongly the applicability of NNT and Padé to realistic problems where the exact function is unknown. The MEM is rather robust to external noise, but for high precision data its quality seems inferior to both Mishchenko's method and Padé.

C. Band gap model

The final model function is useful to investigate realistic features that are not close to zero energy. We again use a non-interacting Hubbard model on a square lattice with nearest-neighbor hopping. Differently from the previous test, we now consider two non-degenerate bands. Again, we take $4t = 1$ for the half bandwidth. The two bands are shifted with respect to each other by an energy $E_{\text{shift}} = 3$. Since the half bandwidth of both bands is 1, there is an indirect gap of $E_{\text{gap}} = 1$. In Equation (11), this would correspond to two bands with shifted disper-

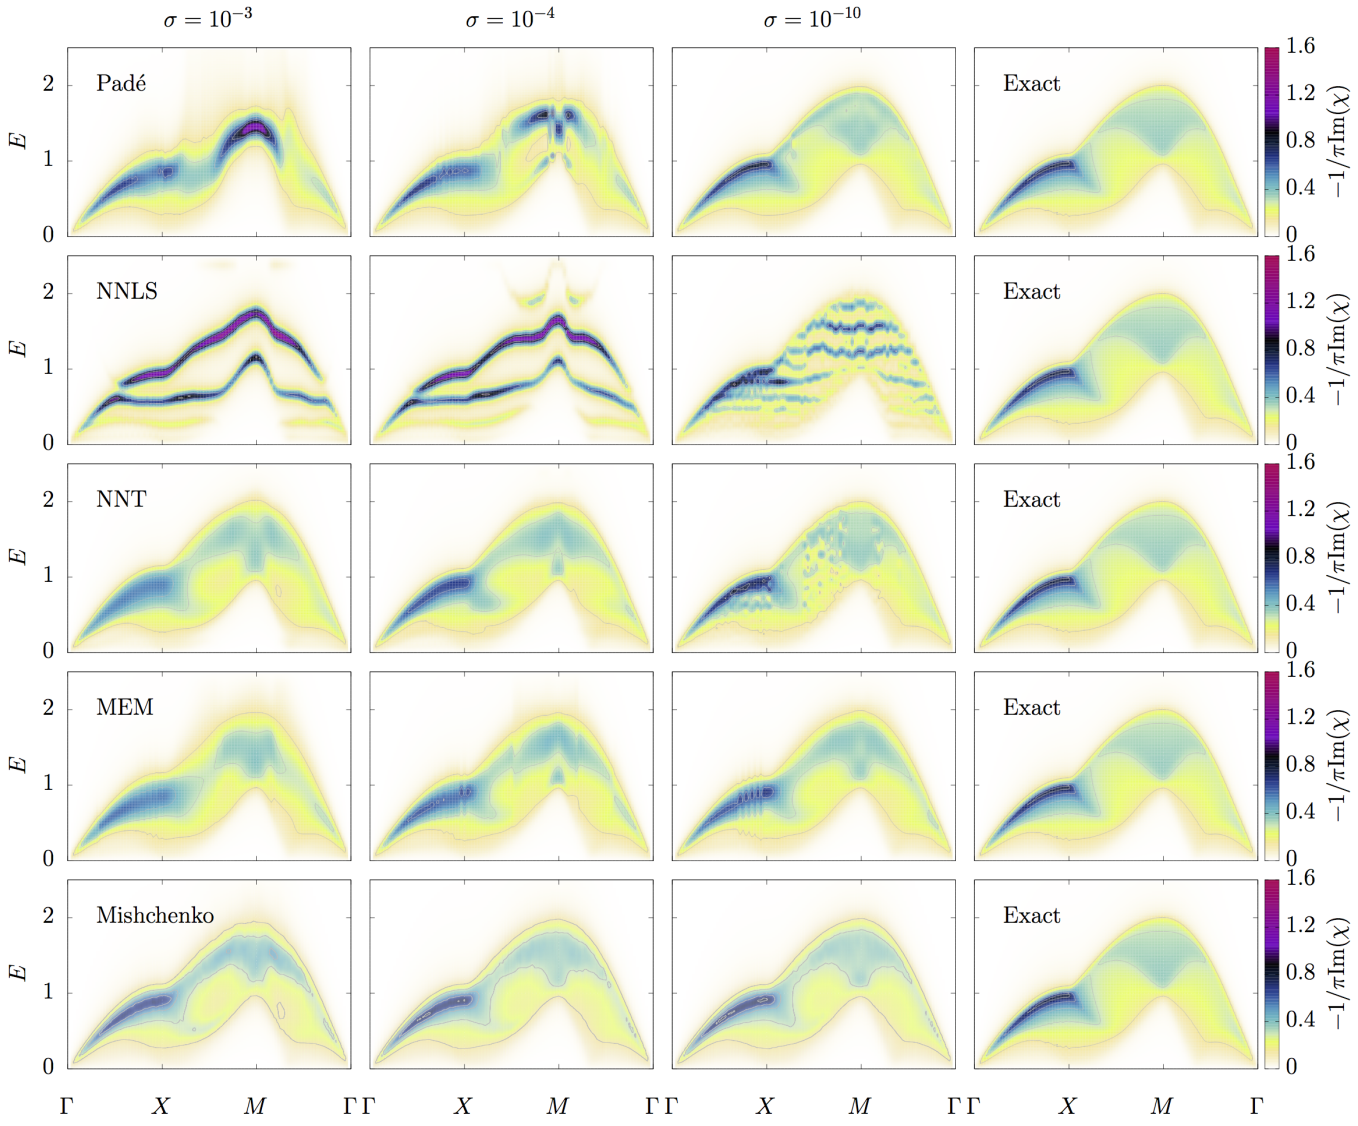

Figure 4. RPA spectrum $(-\frac{1}{\pi}\mathrm{Im}[\chi(\omega + i\delta)])$ of the two-dimensional doped Hubbard model along the Brillouin zone path $\Gamma \to X \to M \to \Gamma$, obtained using Padé, NNLS, NNT, MEM and Mishchenko's methods. The noise level varies among the values $\sigma = 10^{-3}$, $\sigma = 10^{-4}$ and $\sigma = 10^{-10}$. All panels on the right hand side show the exact spectrum, for an easier comparison. The color map denotes the intensity of the spectra and contour lines are added at $-1/\pi\mathrm{Im}(\chi) = 0, 0.15, 0.30, 0.45, 0.60, 0.75$. For the NNLS plots not all contour lines are displayed (see main text). The spectra for the M-point separately are also reported in Fig. 5.

sions $\epsilon_{\mathbf{p}}$ and $\epsilon'_{\mathbf{p}} = \epsilon_{\mathbf{p}} + 3$. We consider both excitations within one band as well as between the two bands. The chemical potential $\mu = 1.7$ lies in the gap between the two bands. This model could correspond to the presence of a magnetic field that splits the dispersion of the electrons with spins up and down. The electronic states with spin down have energy $E_\downarrow \in [-1, 1]$, i.e. they are all below the Fermi energy. The electronic states with spin up, on the other hand, have energy $E_\uparrow \in [2, 4]$. i.e. they are all above the Fermi energy. Excitations between the two bands have energies between $E = 1$ and $E = 5$.

In Fig. 6 the analytical continuation of the momentum resolved RPA susceptibility for all methods is reported. As above, the panels in the first three columns correspond to data for various level of noise σ, while the panels in the last column illustrate the exact spectrum on the real axis. At the Γ-point, the exact spectrum shows the presence of a single peak at $E = 3$. This essential feature is captured by all methods and for all levels of noise, although the MEM and the NNT method exhibit the usual tendency to broaden the peak. Going from Γ to X, the single peak splits into two distinct peaks, with some spectral weight remaining in the intermediate region. To illustrate this, Fig. 7(a) shows the spectrum

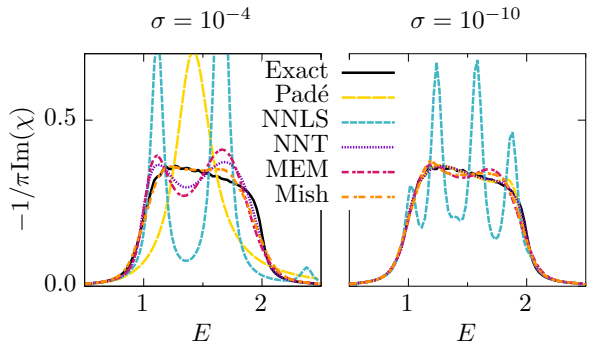

Figure 5. Profile at the M-point of the RPA spectrum of the doped Hubbard model reported in Fig. 4. Data for two noise levels $\sigma = 10^{-4}$ and $\sigma = 10^{-10}$ are shown.

half way between Γ and X. At $\sigma = 10^{-4}$ MEM, NNT and Mishchenko's method all capture the widening of the peak as we move away from the Γ-point. However only Mishchenko's method captures the gradual splitting into two peaks at this point. MEM and NNT find the splitting closer to X. For smaller noise, i.e. for $\sigma = 10^{-10}$, the MEM is also capable of resolving the two-peak structure at $\frac{1}{2}(\Gamma \to X)$. NNLS, NNT and Padé fail to offer an acceptable description of the spectrum at this point, since they also introduce one or more spurious peaks in between the two physical peaks.

At the M point, the exact spectrum has a sharp peak at $E = 3$ with a very broad base. All methods, except for NNLS, give the correct width, but are not very accurate on the symmetric character of the spectrum. This improves for lower noise levels, but the sharp peak in the middle stays slightly broadened, as can be seen in Fig. 7(b). These data show that Mishchenko's method is the only one to offer a fully satisfactory description of the spectrum for $\sigma = 10^{-4}$. It is also worth noticing that although the Padé method does not fail too badly at M, the general trend for $\sigma = 10^{-4}$ observed in Fig. 6 is rather negative and violates symmetry constraints. For even smaller noise, $\sigma = 10^{-10}$, all methods but NNLS give a very good description of the spectrum at the M-point.

In summary, for this model function, Mishchenko's method is stable with respect to input noise and clearly performs best of all methods for $\sigma = 10^{-3}$. However, great care has to be taken concerning the number of updates in the sampling chain, as is investigated in Appendix B. The Padé method performs quite well for high precision data but fails badly when noise is present. Finally, for intermediate levels of noise, MEM and NNT also lead to reasonable results, although Mishchenko's method is still superior.

D. Computational time

As a final note, we want to make a brief comment on the computational effort required to perform the analytical continuation of these functions. Momentum resolved calculations, such as the RPA spectra for the Hubbard model shown above, typically involve linear dimensions of approximately 100 sites. Even considering a minimal effort, analytical continuation is required on a path through the high symmetry points of the lattice, involving about a hundred analytical continuations. Additionally, the analytical continuation should be done several times with different continuation parameters to ascertain its stability. Therefore, the computational time required to perform a single continuation becomes a factor of practical importance.

It is not simple to compare the fundamental computational effort required by the five methods investigated here, due that the employed codes have different degrees of optimization and parallelization. Nevertheless, we can still provide some estimates. The fastest methods are undoubtedly NNLS and NNT, which require less than a second per continuation. The MEM and Padé are slower, and usually require from several to a hundred seconds per continuation, depending on the parameters used. Mishchenko's method is by far the most demanding method among those tested here. To obtain properly converged results one has to perform the stochastic sampling for a time ranging from a few hours to weeks. The effective time can be significantly reduced by using a parallelized code, like we do in the present work. However, parallelization is easy to implement only over global updates in the Monte Carlo chain, while local updates have still to be handled serially. This limits the scalability of the code, and therefore the reduction in computational time, to about hundred CPUs. Even in this case, we must notice that it is not always possible to dedicate such a large amount of resources to the analytical continuation problem. This sets a limit on the applicability of Mishchenko's method in current research. We are currently exploring an alternative formulation of Mishchenko's method based on a graphical processing unit (GPU) [46].

VI. CONCLUSIONS

We have presented results for the analytical continuation of bosonic functions by five different methods and using three realistic test cases, corresponding to different scenarios of physical relevance for strongly correlated systems. An overview of the major features emphasized in the previous section is given in Table I. The first conclusion of our work is that, at the moment, for low input accuracy, none of the existing methods is good enough to be able to continue all different types of two-particle quantities from the Matsubara frequencies to real energies. Nevertheless, a combination of several methods can

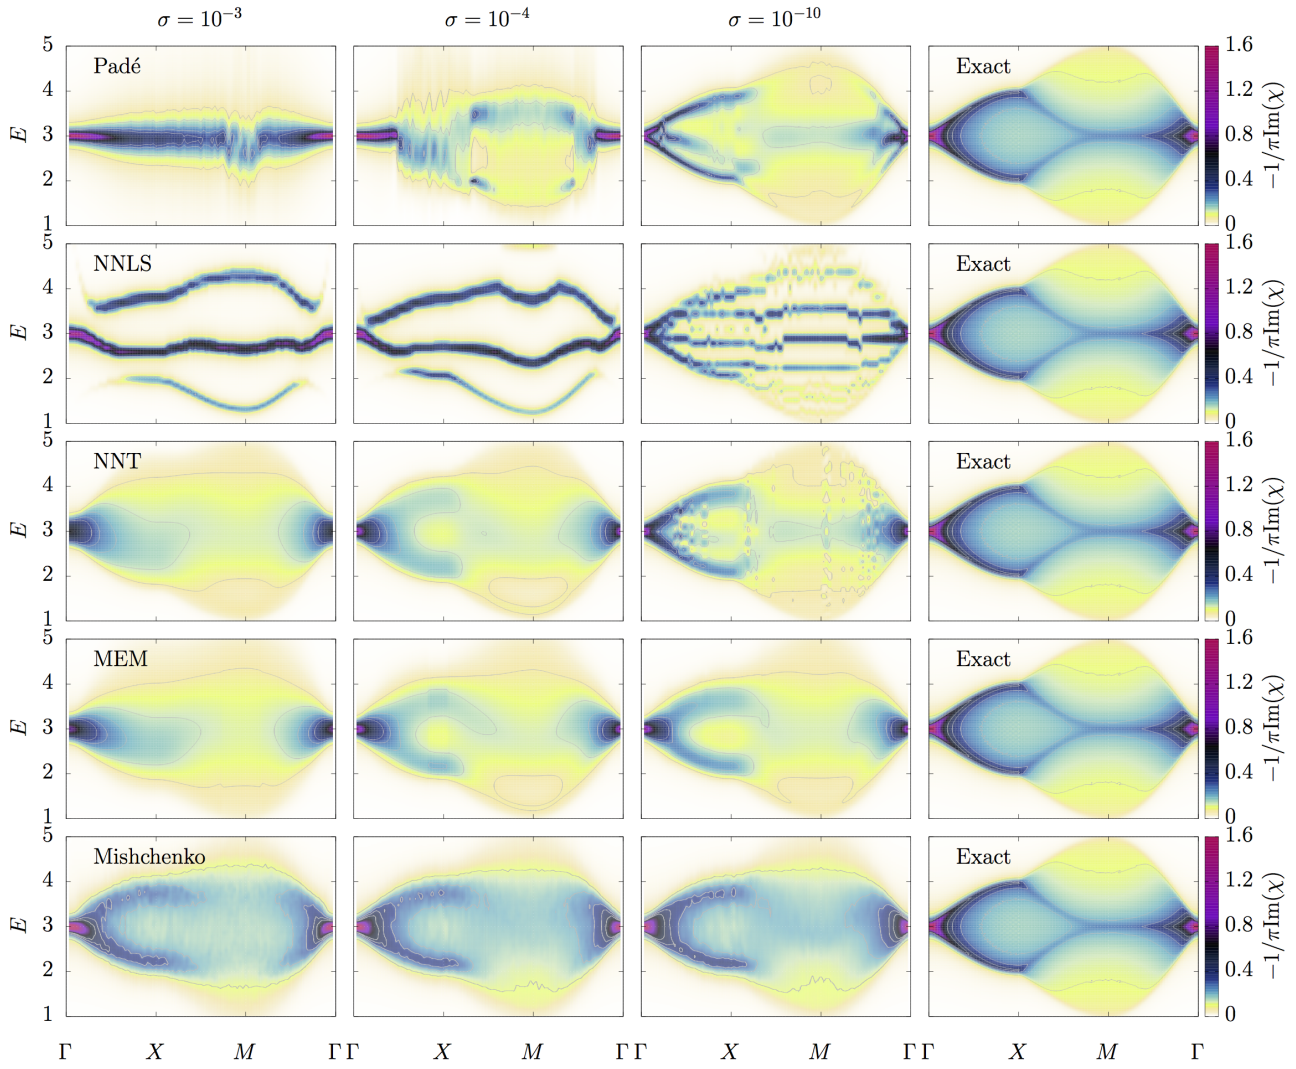

Figure 6. RPA spectrum of the band gap model ($-\frac{1}{\pi}\mathrm{Im}[\chi(\omega + i\delta)]$) along the Brillouin zone path $\Gamma \to X \to M \to \Gamma$, obtained using Padé, NNLS, NNT, MEM and Mishchenko's method. The noise level varies among the values $\sigma = 10^{-3}$, $\sigma = 10^{-4}$ and $\sigma = 10^{-10}$. All right panels show the exact spectrum. The color map denotes the intensity of the spectra and contour lines are added at $-1/\pi\mathrm{Im}(\chi) = 0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6$. The spectra at the M-point and at $\frac{1}{2}(\Gamma \to X)$ poit are reported separately in Fig. 7.

give enough information to reconstruct the physical picture behind the analytical continuation problem. Table I can offer useful guidelines to perform this type of analysis.

Conversely, for high precision input data, our results are encouraging. Among all methods, the Padé averaging scheme [44] is the one that performs best when the numerical noise is absent. For complicated spectra, such as that of the band gap model, even Padé cannot completely resolve the spectrum for $\sigma = 10^{-10}$, and it is indeed not even clear that it performs better than Mishchenko's method, due to small violations of the symmetry characterizing the exact spectrum. In general, Mishchenko's method seems to be the most robust approach across various levels of noise. As a matter of fact, its results do not seem very affected by the different noise levels for all the tested functions. The accuracy of Mishchenko's method is also good, although care is needed when choosing the numerical parameters of the stochastic sampling. For optimal parameter values, the accuracy of Mishchenko's method is definitely good.

Overall, a single sharp peak close to zero energy is resolved by all methods even for a relatively high noise. The MEM, NNT method and Mishchenko's method are able to resolve broad features. However, these methods also have the tendency to smear several pronounced peaks, sometimes into a single broad feature. In Mishchenko's method, this problem is related to the

Table I. Summary of the test performance.

	σ	Padé	NNLS	NNT	MEM	Mishchenko
2 close poles	10^{-2}	1 peak	Broadened, positions wrong	Broadened, positions wrong	1 peak	Broadened, positions wrong
	10^{-4}	Perfect	Good	Good	Good	Good
2 separated poles	10^{-2}	1 peak	Good	Good, broadened	1 peak	Good, broadened
	10^{-4}	Perfect	Good	Good	Good, broadened	Perfect
1 sharp feature	10^{-3}	Good	Spurious features	Broadened	Broadened	Good
	10^{-10}	Excellent	Spurious features	Broadened, unstable	Broadened	Good
Broad, flat feature	10^{-3}	Too sharp	2 sharp peaks	Good, unphysical minima	Good, unphysical minima	Good
	10^{-10}	Good	Spurious features	Good	Good	Good
2 peaks in broad plateau	10^{-3}	1 peak	Wrong positions +additional peak	1 broad peak	1 broad peak	Good, broadened
	10^{-10}	Additional peak	Spurious peaks	Spurious peaks	Good, broadened	Good, broadened
1 peak in broad plateau	10^{-3}	1 peak	Spurious peaks	Plateau	Broad peak +plateau	Chopped triangle
	10^{-10}	Good	Spurious peaks	Good	Good	Good

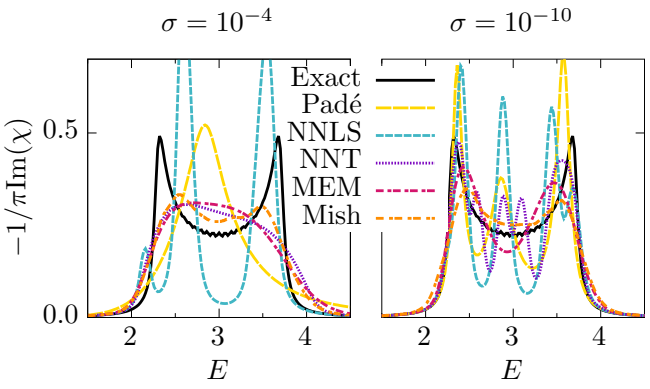

(a)Profile at $\frac{1}{2}(\Gamma \to X)$-point

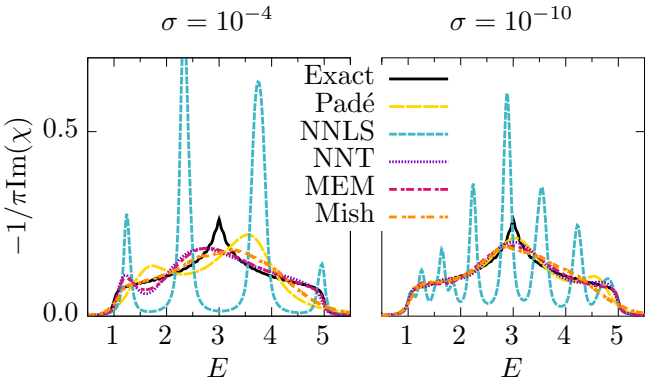

(b)Profile at M-point

Figure 7. Profile at the $\frac{1}{2}(\Gamma \to X)$ and M-point of the RPA spectrum of the band gap model in Fig. 6 for two noise levels $\sigma = 10^{-4}$ and $\sigma = 10^{-10}$.

number of local updates performed in the sampling. As shown in Appendix B, increasing the number of local updates leads to more peaked spectra, but those are not always in better agreement with the exact results. In general, the tendency to smear pronounced features into broad peaks and, at the same time, to create spurious peaks makes it difficult to know if a calculated broad . But on the other hand, sometimes spurious structure is created. This makes it difficult to know if a calculated broad spectral feature really is a broad feature or a combination of several peaks. As an example of this, the MEM continuation of the two-pole function for $E_2 = 1.2$ at $\sigma = 10^{-4}$ (Fig. 2) looks similar to the MEM continuation of the two-dimensional doped Hubbard model at the M-point for the same noise level (Fig. 5), although the exact functions are very different.

The very broad modes that appear, e.g., around the M point of the doped two-dimensional Hubbard model, are difficult to resolve correctly. Functions with two distinct poles, on the other hand, are much easier. NNLS, in particular, performs well for the two-pole model whereas it performs very poorly for more complicated spectra. This shows the importance of varied, realistic test functions for assessing the potential of each continuation method.

Finally, we compare our work to previous studies. Two previous works [12, 27] found that Padé generally performs poorer than the other methods. Huang *et al.* [12] found that MEM outperforms Padé for a spectrum generated from a combination of two Gaussian distributions. In particular, their results show that the Padé method sometimes merges two peaks, just as we have seen in the two-pole spectra. In Ref. [27], the authors test several methods of analytical continuation on the optical con-

ductivity for relatively high noise levels ($\sigma = 10^{-2}$ to 10^{-3}). They find that all methods (Padé, Singular Value Decomposition, sampling and MEM) perform similarly, but that the Padé method generally gives slightly less accurate results and sometimes finds unphysical continuations. However, in both studies, the averaging scheme for Padé [44] was not used. This averaging scheme solves the problem of unphysical continuations and should improve the overall performance of the Padé method. We would also like to stress two more fundamental conclusions of this study. First, we want to emphasize the importance of varied, realistic test functions in the assessment of continuation algorithms. The dependence of the performance of each method on the type of function investigated is a clear indication that no analysis can be meaningful without exploring a proper number of test cases. Second, we want to stress that our findings fully support what was already claimed in Ref. [27] , namely, despite no universal tool currently available, a better insight into the continuation problem can be obtained by using a combined approach of several methods.

VII. ACKNOWLEDGMENTS

This work was sponsored by the Swedish Research Council (VR), the Swedish strategic research programme eSSENCE and the Knut and Alice Wallenberg foundation (KAW). E.G.C.P.v.L. and M.I.K acknowledge support from ERC Advanced Grant 338957 FEMTO/NANO. The authors acknowledge the computational resources provided by the Swedish National Infrastructure for Computing (SNIC) at the PDC Center for High Performance Computing at the KTH Royal Institute of Technology, at Uppsala Multidisciplinary Center for Advanced Computational Science (UPPMAX) and at Chalmers Centre for Computational Science and Engineering (C3SE).

Appendix A: Imposing mirror symmetry

The MEM, the NNT and the NNLS method conserve the symmetry of the spectra $\rho(E) = -\rho(-E)$, since they use the symmetric definition of the Hilbert transform, Eq.(2). The Padé method does not have such symmetry imposed. For fermionic one-particle Greens functions and self-energies with zero imaginary part at the origin of the complex plane, it is shown [44] that the continuation improves by taking a few negative Matsubara frequencies into account to impose the mirror symmetry. Such a procedure, with a single negative Matsubara frequency, has also been used for the optical conductivity [27]. We have tried the same scheme for the Padé continuation of the bosonic functions in this paper, but in most cases imposing mirror symmetry actually gave a worse continuation. In particular, when the exact spectrum has a broad, rather flat spectrum, such as at the M-point

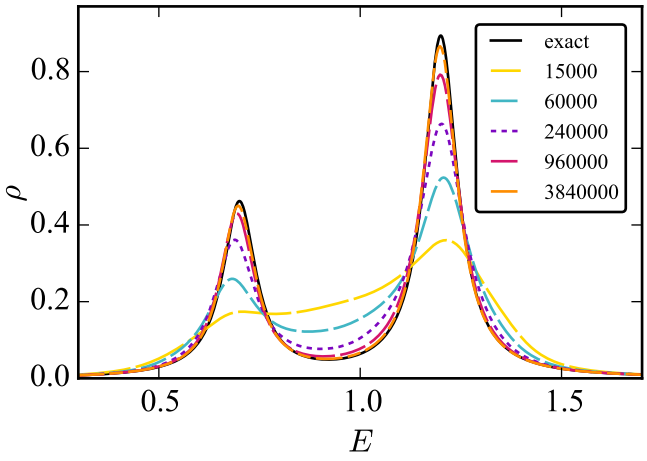

(a)The two-pole function in the presence of input noise with $\sigma = 10^{-4}$

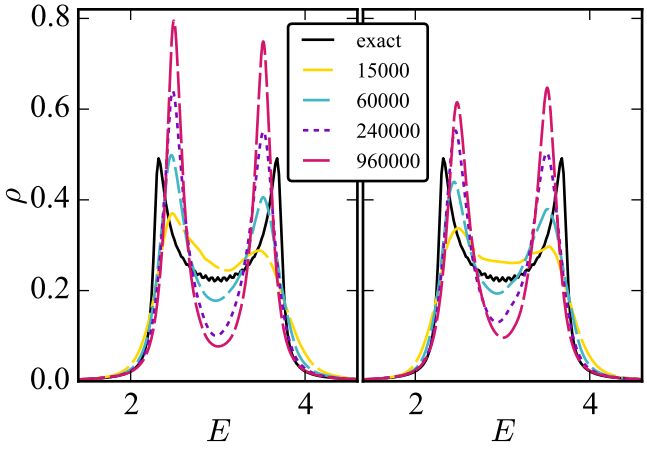

(b)The band gap model at k-point $\frac{1}{2}(\Gamma \to X)$, for two different noise levels of $\sigma = 10^{-3}$ (left) and $\sigma = 10^{-4}$ (right).

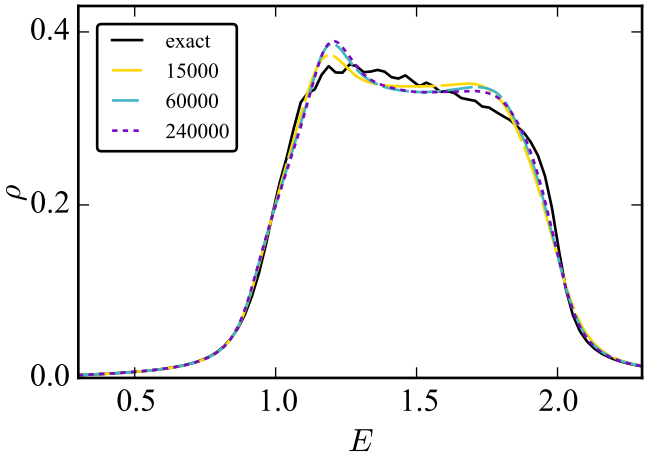

(c)The doped Hubbard model at the M-point and Matsubara noise level $\sigma = 10^{-4}$.

Figure 8. The dependence on the number of iterations (see legend) performed with the Mishchenko method.

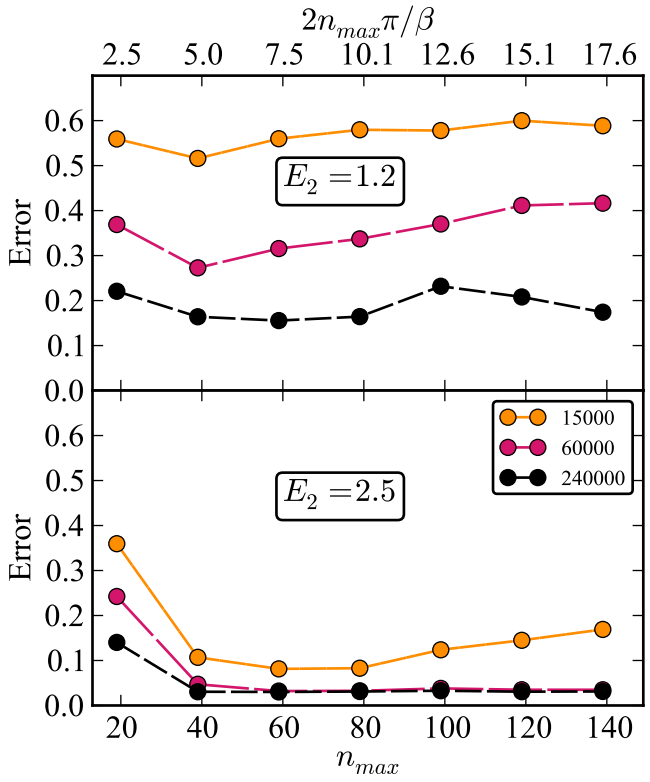

Figure 9. Integrated real axis error for the two-pole model as a function of the highest Matsubara point index $n_{\max}$ used in Mishchenko's method. Different number of local iterations are investigated and are indicated by the legend. Upper panel $a_1 = 0.1$, $a_2 = 0.335663$, $E_1 = 0.7$ and $E_2 = 1.2$, lower panel $a_1 = 0.1$, $a_2 = 0.7$, $E_1 = 0.7$ and $E_2 = 2.5$, as defined in Eq. 9. Matsubara noise level $\sigma = 10^{-4}$ is used.

of the doped Hubbard model, imposing the symmetry can lead to spurious sharp features. Broad modes in the spectrum are related to branch cuts in the complex function $\chi(z)$. The Padé approximant uses a finite number of poles to simulate $\chi(z)$. When the mirror symmetry is imposed on the Padé approximant, the poles move closer to the real axis, which results in sharper spectral features. Hence broad features are harder to accurately simulate by imposing mirror symmetry.

Appendix B: Convergence of Mishchenko's sampling method

The stochastic sampling method by Mishchenko *et al.* [40] consists of many iterative steps. In every step, the configuration of the rectangles is altered. In the main text results with 15000 iterations are presented, which is a suitable setup to obtain reasonable results in a decent computational time (30 cpu hours using 128 cpus). However, in principle the method should be converged carefully with respect to the number of iterations. In Fig. 8 we increase that number by a factor of 4 several times, and show the resulting spectral function. For the two-pole model (see Fig. 2) it is clear that increasing the number of iterations improves the resulting spectrum, see Fig. 8(a). However for a more complicated spectrum, such as the band gap model at the k-point $\frac{1}{2}(\Gamma \to X)$ (see Fig. 7(a)), convergence with respect to the number of iterations does not lead to an improvement, see Fig. 8(b). This issue is particularly pronounced when the input noise is big. The tendency to produce too sharp features resembles the overfitting occurring in the NNLS method. For the doped Hubbard model at the M-point (see left panel in Fig. 5) the spectra converge already for 15000 iterations, see Fig. 8(c). In summary, the optimal number of iterations is highly dependent on the function under investigation. With increasing number of iterations, the computational cost increases as well as the risk of overfitting to the input Matsubara data. Hence care should be taken when choosing the number of iterations.

Further, It is interesting to repeat the analysis on the accuracy of the continuation with respect to the number of Matsubara points (Fig. 3) while varying the number of local updates. These results are illustrated in Fig. 9. While increasing the number of local updates leads to a better continuation, the general trend reported in Fig. 9 is very similar to Fig. 3. Increasing the number of Matsubara frequencies the quality of the continuation reaches saturation very quickly and with a threshold that does not seem to depend on the number of local updates. A posteriori, this shows that the approach used in Fig. 3, i.e. keeping fixed the number of local updates while increasing the number of Matsubara points, is justifiable.

[1] V. Anisimov and Y. Izyumov, *Electronic Structure of Strongly Correlated Materials*, edited by M. Cardona, Springer Series in Solid-State Sciences (Springer-Verlag, Berlin Heidelberg, 2010).

[2] W. Metzner and D. Vollhardt, Phys. Rev. Lett. **62**, 324 (1989).

[3] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, Rev. Mod. Phys. **68**, 13 (1996).

[4] A. I. Lichtenstein and M. I. Katsnelson, Phys. Rev. B **57**, 6884 (1998).

[5] G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet, and C. A. Marianetti, Rev. Mod. Phys. **78**, 865 (2006).

[6] J. W. Negele and H. Orland, *Quantum Many-Particle Systems* (Addison-Wesley Publishing Company, 1988).

[7] G. D. Mahan, *Many Particle Physics, Third Edition* (Plenum, New York, 2000).

[8] L. Boehnke and F. Lechermann, Phys. Rev. B **85**, 115128 (2012).

[9] T. Ayral, P. Werner, and S. Biermann, Phys. Rev. Lett. **109**, 226401 (2012).

[10] P. Hansmann, T. Ayral, L. Vaugier, P. Werner, and S. Biermann, Phys. Rev. Lett. **110**, 166401 (2013).

[11] T. Ayral, S. Biermann, and P. Werner, Phys. Rev. B **87**, 125149 (2013).

[12] L. Huang, T. Ayral, S. Biermann, and P. Werner, Phys. Rev. B **90**, 195114 (2014).

[13] E. G. C. P. van Loon, H. Hafermann, A. I. Lichtenstein, A. N. Rubtsov, and M. I. Katsnelson, Phys. Rev. Lett. **113**, 246407 (2014).

[14] H. Hafermann, E. G. C. P. van Loon, M. I. Katsnelson, A. I. Lichtenstein, and O. Parcollet, Phys. Rev. B **90**, 235105 (2014).

[15] A. Galler, C. Taranto, M. Wallerberger, M. Kaltak, G. Kresse, G. Sangiovanni, A. Toschi, and K. Held, Phys. Rev. B **92**, 205132 (2015).

[16] A. S. Mishchenko, N. Nagaosa, G. De Filippis, A. de Candia, and V. Cataudella, Phys. Rev. Lett. **114**, 146401 (2015).

[17] A. A. Katanin, A. I. Poteryaev, A. V. Efremov, A. O. Shorikov, S. L. Skornyakov, M. A. Korotin, and V. I. Anisimov, Phys. Rev. B **81**, 045117 (2010).

[18] L. Boehnke, H. Hafermann, M. Ferrero, F. Lechermann, and O. Parcollet, Phys. Rev. B **84**, 075145 (2011).

[19] Q. Si and J. L. Smith, Phys. Rev. Lett. **77**, 3391 (1996).

[20] H. Kajueter, Ph.D. thesis, Rutgers University (1996).

[21] J. L. Smith and Q. Si, Phys. Rev. B **61**, 5184 (2000).

[22] R. Chitra and G. Kotliar, Phys. Rev. Lett. **84**, 3678 (2000).

[23] R. Chitra and G. Kotliar, Phys. Rev. B **63**, 115110 (2001).

[24] A. N. Rubtsov, M. I. Katsnelson, and A. I. Lichtenstein, Annals of Physics **327**, 1320 (2012).

[25] E. G. C. P. van Loon, A. I. Lichtenstein, M. I. Katsnelson, O. Parcollet, and H. Hafermann, Phys. Rev. B **90**, 235135 (2014).

[26] E. G. C. P. van Loon, M. I. Katsnelson, and M. Lemeshko, Phys. Rev. B **92**, 081106 (2015).

[27] O. Gunnarsson, M. W. Haverkort, and G. Sangiovanni, Phys. Rev. B **82**, 165125 (2010).

[28] O. Gunnarsson, M. W. Haverkort, and G. Sangiovanni, Phys. Rev. B **81**, 155107 (2010).

[29] M. Casula, A. Rubtsov, and S. Biermann, Phys. Rev. B **85**, 035115 (2012).

[30] R. N. Silver, D. S. Sivia, and J. E. Gubernatis, Phys. Rev. B **41**, 2380 (1990).

[31] M. Jarrell and J. Gubernatis, Physics Reports **269**, 133 (1996).

[32] R. Bryan, European Biophysics Journal **18**, 165 (1990).

[33] J. E. Gubernatis, M. Jarrell, R. N. Silver, and D. S. Sivia, Phys. Rev. B **44**, 6011 (1991).

[34] S. Fuchs, T. Pruschke, and M. Jarrell, Phys. Rev. E **81**, 056701 (2010).

[35] A. Dirks, P. Werner, M. Jarrell, and T. Pruschke, Phys. Rev. E **82**, 026701 (2010).

[36] C. L. Lawson and R. J. Hanson, *Solving Least Squares Problems* (Society for Industrial and Applied Mathematics Philadelphia, Philadelphia, 1995).

[37] P. C. Hansen, in *Computational Inverse Problems in Electrocardiology, ed. P. Johnston, Advances in Computational Bioengineering* (WIT Press, 2000) pp. 119–142.

[38] K. S. D. Beach, R. J. Gooding, and F. Marsiglio, Phys. Rev. B **61**, 5147 (2000).

[39] H. J. Vidberg and J. W. Serene, J. Low Temp. Phys. **29**, 179 (1977).

[40] A. S. Mishchenko, N. V. Prokof'ev, A. Sakamoto, and B. V. Svistunov, Phys. Rev. B **62**, 6317 (2000).

[41] A. S. Mishchenko, in *Verlag des Forschungszentrum Jülich, 2012. - 978-3-89336-796-2*, edited by E. Pavarini, E. Koch, F. Anders, and M. Jarrell (2012).

[42] E. Gull, A. J. Millis, A. I. Lichtenstein, A. N. Rubtsov, M. Troyer, and P. Werner, Rev. Mod. Phys. **83**, 349 (2011).

[43] L. D. Landau and E. M. Lifshitz, "Statistical physics," (1980).

[44] J. Schött, I. L. M. Locht, E. Lundin, O. Grånäs, O. Eriksson, and I. Di Marco, Phys. Rev. B **93**, 075104 (2016).

[45] D. Pines and P. Nozières, *The Theory of Quantum Liquids: Normal Fermi liquids* (W.A. Benjamin, Philadelphia, 1966).

[46] J. Nordström, J. Schött, I. L. M. Locht, and I. Di Marco, "A GPU code for analytic continuation through a sampling method," (2016), unpublished.